

Efficient Calculation of Short-Range Coulomb Energies

ROSS D. ADAMSON,¹ JEREMY P. DOMBROSKI,² PETER M. W. GILL¹

¹Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom

²HortResearch, Private Bag 11030, Palmerston North, New Zealand

Received 5 October 1998; accepted 1 February 1999

ABSTRACT: An efficient algorithm for the calculation of short-range Coulomb energies is examined. The algorithm uses a boxing scheme and a prescreening for negligible integrals to evaluate the short-range Coulomb energy via computational work that scales only linearly with the size of the system.

© 1999 John Wiley & Sons, Inc. J Comput Chem 20: 921–927, 1999

Keywords: CASE; linear scaling; integrals; KWIK; boxing schemes

Introduction

The last few years have witnessed rapid growth in the development of methods that compute the Coulomb energy of a system of N localized distributions in only $O(N)$ operations.^{1–6} Each aims to avoid the evaluation of $O(N^2)$ two-electron integrals that, due to their sheer number, have become the rate-limiting step for self-consistent field calculations. The first $O(N)$ algorithm, the fast multipole method (FMM) of Greengard and Rokhlin,¹ was designed for point charges and achieves its linear scaling by partitioning physical space into boxes. All interactions are then classified as either ‘well-separated’ (WS) or ‘not well-

separated’ and this splits the Coulomb problem into two parts. The not-WS interactions, of which there are $O(N)$ in spatially extended systems, are handled using traditional techniques.^{7,8} The WS interactions are treated in linear work by calculating the multipole moments of groups of particles and then interacting distant groups via their multipoles. This avoids explicit treatment of each particle–particle interaction.

Inspired by this and the more traditional Ewald approaches,^{9–11} Dombroski et al. introduced a second linear scheme,⁵ the KWIK algorithm, wherein the Coulomb operator (rather than physical space) is partitioned into two parts:

$$\frac{1}{r} = \frac{\operatorname{erfc}(\omega r)}{r} + \frac{\operatorname{erf}(\omega r)}{r}, \quad (1)$$

one singular but short-ranged, the other nonsingular but slowly decaying. As in the FMM this splits the Coulomb problem into two parts, each of which

Correspondence to: Dr. R. D. Adamson; e-mail: ross@theor.ch.cam.ac.uk

Contract/grant sponsors: Cambridge Commonwealth Trust; Trinity College; ORS Awards Scheme; Q-Chem, Inc.

can be treated separately. The long-range operator is very flat and we showed^{12–15} that even its complete neglect has little effect on many chemical properties, provided that the attenuation parameter ω is not too large. Systematic neglect of the long-range energy is called the CASE approximation. Using $\omega = 0.1a_0^{-1}$, CASE typically reproduces atomization energies to within 3 kJ/mol and has almost no effect on wave functions.¹² However, if the attenuation is increased to $\omega = 0.5a_0^{-1}$, the loss in accuracy becomes unacceptably large and a correction for the long-range energy is required for reliable reproduction of most chemical properties.¹⁵

All CASE results published thus far^{12–15} were obtained using an $O(N^2)$ implementation that does not exploit the rapid decay of the erfc function. This article describes several algorithms to calculate short-range Coulomb energies efficiently and, in particular, presents a fully linear CASE implementation. This should be useful both in its own right and as a necessary requirement for $O(N)$ calculations using any of the KWIK family of methods^{5,16–19} where a correction is made for the neglected term. We implemented our approach within the Q-Chem 1.2 program.²⁰

Integral Generation

PRISM algorithms^{7,8} generate the two-electron repulsion integral (ERI),

$$(\mathbf{ab}|\mathbf{cd}) = \iint \frac{\phi_a(\mathbf{r}_1)\phi_b(\mathbf{r}_1)\phi_c(\mathbf{r}_2)\phi_d(\mathbf{r}_2)}{r_{12}} d\mathbf{r}_1 d\mathbf{r}_2, \quad (2)$$

over contracted Gaussian-type basis functions,

$$\phi_a(\mathbf{r}) = (x - A_x)^{a_x}(y - A_y)^{a_y}(z - A_z)^{a_z} \times \sum_{i=1}^{K_a} D_{ai} e^{-\alpha_i |\mathbf{r} - \mathbf{A}|^2}, \quad (3)$$

by applying recurrence relations to the fundamental $[m]$ integrals

$$[m] = U(2\theta^2)^{m+1/2} G_m(T), \quad (4)$$

$$G_m(T) = (2/\pi)^{1/2} \int_0^1 t^{2m} \exp(-Tt^2) dt, \quad (5)$$

$$T = \theta^2 R^2, \quad (6)$$

$$\theta^2 = 1 / \left[\frac{1}{\zeta} + \frac{1}{\eta} \right] \quad (7)$$

and contracting. All symbols have their standard meanings.⁷ If T is less than a critical value T_{crit} , $G_m(T)$ and the function $\exp(-T)$ are calculated via interpolation. If T is greater than T_{crit} , the distributions $\phi_a\phi_b$ and $\phi_c\phi_d$ overlap negligibly and (4) reduces to the classical multipole formula²¹

$$[m] = \left(\frac{U}{R} \right) \left(\frac{1}{R^2} \right) \left(\frac{3}{R^2} \right) \cdots \left(\frac{2m-1}{R^2} \right), \quad (8)$$

which can be computed efficiently via recursion. For a full review of integral theory see ref. 7.

In CASE calculations, the $1/r_{12}$ operator is replaced by the attenuated Coulomb operator

$$f(r_{12}) = \text{erfc}(\omega r_{12})/r_{12}, \quad (9)$$

and we therefore require

$$= (\mathbf{ab}|\mathbf{cd})_\omega = \iint \phi_a(\mathbf{r}_1)\phi_b(\mathbf{r}_1) \frac{\text{erfc}(\omega r_{12})}{r_{12}} \times \phi_c(\mathbf{r}_2)\phi_d(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2. \quad (10)$$

As before, these can be generated from the fundamental integrals

$$[m]_\omega = U \left\{ (2\theta^2)^{m+1/2} G_m(T) - (2\theta_\omega^2)^{m+1/2} G_m(T_\omega) \right\}, \quad (11)$$

$$T_\omega = \theta_\omega^2 R^2, \quad (12)$$

$$\theta_\omega^2 = 1 / \left[\frac{1}{\zeta} + \frac{1}{\eta} + \frac{1}{\omega^2} \right], \quad (13)$$

which is the same as the traditional $[m]$ except for a correction term dependent on the attenuation parameter ω . Thus, we need change only one subroutine in Q-Chem to allow calculation of the CASE energy. Note, however, that such a change does not take advantage of the rapid decay of this short-range operator and will therefore still scale as $O(N^2)$ with the size of the system.

Our algorithm follows the standard PRISM method,^{7,8} except that two new variables ($T_\omega < T$ and $\theta_\omega^2 < \theta^2$) are calculated and an extra G_m interpolation is required. Because of this, the calculation of CASE integrals (11) takes slightly longer than standard integrals (4).

However, it should be noted that if $T_\omega > T_{\text{crit}}$, the distributions $\phi_a\phi_b$ and $\phi_c\phi_d$ are sufficiently distant that their interaction over the attenuated operator (9) is negligible. In this event, $[m]_\omega$ vanishes rather than reducing to the multipole expansion (8). The integral screening presented below relies on this property.

Boxing Scheme

A short-range operator decays much faster than the Coulomb operator. Consequently, beyond a certain size, the number of nonnegligible integrals in a system grows only as $O(N)$. We therefore seek an $O(N)$ algorithm for identifying these non-negligible interactions. The simplest way to do this is to divide the system into cubical boxes (or some other space-filling shape²²) and then interact the contents of each box only with neighboring boxes. There are a variety of different boxing schemes that one can apply.^{2,23,24} However, we chose to adopt the scheme used for the near-field continuous FMM. The additional complexity of alternative methods is unlikely to offer significant benefit and the CFMM implementation in Q-Chem is already highly optimized.

The first step is to determine the number of boxes in each Cartesian direction. The side length of each box is taken to be the minimum distance over which the short-range operator is negligible; for CASE integrals this is the value of r when

$$\frac{\text{erfc}(\omega r)}{r} = \text{accuracy required.} \quad (14)$$

Given ω and the integral accuracy required, it can be found using the Newton–Raphson method.²⁵ Table I lists some typical values for the side length of a box. The number of boxes in a direction is then the length of the molecule in that direction, divided by the side length.

Following the CFMM, our algorithm assigns to each shell pair a box and a WS index. The latter is the number of box lengths over which the shell pair's charge is nonnegligible.²⁶ A contracted shell pair, $\phi_a\phi_b$, contains several primitives and sometimes these may be centered in different boxes. If

this is the case, the contracted shell pair is split into two (or more) shell pairs, one for each box that the original shell pair centers straddle.

The boxes are then looped over in the same way as in the near-field CFMM algorithm: each shell pair of a particular box is interacted with the shell pairs of boxes within $WS + 1$ multiples of the box length.

Integral Screening

Although the above boxing scheme is $O(N)$, it is still inefficient because all interactions within a cube the size of $WS + 1.5$ times the box length, centered at the current box center, are treated. Yet, only those within a sphere of radius $(WS + 1)$ times the point charge interaction distance (PID), centered at the current point, will be nonnegligible. If we define the efficiency to be

$$\varepsilon = \frac{\text{no. interactions that should be computed}}{\text{no. interactions that are computed}}, \quad (15)$$

then the efficiency of the boxing scheme above is only 15%! This is an inherent deficiency in using cubes as our space-filling shape. The efficiency can be improved by shrinking the box size and correspondingly increasing WS values, but this markedly reduces loop lengths in the computer program.

A useful compromise is to introduce Johnson's²⁷ prescreening routine. Just before looping over all the shell pairs to be interacted with a given shell pair, we determine the T_ω values for all primitive shell quartets and only if one is less than T_{crit} is the contracted shell pair included. This adds about 15% to the work required in calculating T_ω , but greatly reduces all subsequent integral work.

TABLE I. Some Typical Values for the Box Sidelength (Å)

ω	Required Integral Accuracy				
	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}
0.10	13.67	15.63	17.41	19.04	20.56
0.25	5.79	6.54	7.23	7.86	8.46
0.50	3.01	3.38	3.71	4.02	4.31
1.00	1.56	1.74	1.90	2.06	2.20
2.00	0.81	0.90	0.98	1.05	1.12

Results and Discussion

How many integrals does the CASE approximation remove? Figure 1 plots the number of significant ($> 10^{-6}$) integrals against ω for a linear alkane ($C_{90}H_{182}$), a graphitic sheet ($C_{54}H_{18}$), and a terminated diamond chunk ($C_{84}H_{64}$). The 6-31G* basis is used in all cases. The three molecules may be characterized as 1-dimensional, 2-dimensional, and 3-dimensional, respectively. The removal of integrals is very efficient in the alkane and, beyond

mild attenuation ($\omega = 0.2a_0^{-1}$), most of the integrals are insignificant. The graphene is more compact than the alkane and, as a result, much larger attenuation is required before the pruning begins to take effect. Once the pruning has begun, however, there is a rapid removal of integrals. The effect of CASE on the diamond chunk follows a similar pattern. $C_{84}H_{64}$ is very compact and a significant reduction in the number of integrals does not begin until large ω values. For much larger systems there will be a significant reduction in the number of integrals for modest attenuation. The slope of the tails of these curves is largest for the 3-dimensional system and smallest for the alkane. This is as expected, because increasing ω decreases the PID, which in a 1-dimensional system produces a corresponding decrease in the number of integrals. However, in a 3-dimensional




FIGURE 1. Number $n(\omega)$ of significant ERIs as a function of ω .

system the same decrease in PID removes many more integrals because each $\phi_a\phi_b$ has many more neighbors. Clearly, the efficiency of CASE depends on the shape of the system, as well as the number of atoms present.

Figure 2 presents CPU times required to compute the integrals required for an HF/6-31G* calculation on a series of linear alkanes using $\omega = 0.25a_0^{-1}$ and an integral cutoff of 10^{-9} . Because we are attenuating individual integrals, this approach is applicable to both Coulomb and exchange. The calculations were performed on a single R8000 processor of an SGI Power Challenge. The solid curve shows the quadratic behavior of the conventional $O(N^2)$ algorithm. Just above it lies the curve for a CASE calculation, clearly showing the cost of the extra interpolation for $G_m(T_\omega)$. The poor performance of QCTC/ONX²⁸ is due to the high integral accuracy used. As shown later in this section, ONX becomes more competitive at low integral accuracy. The two near-linear curves are for CASE + screening and CASE + boxing + screening. The remaining curve is for CASE + boxing. It is perhaps surprising that integral screening, although formally $O(N^2)$, performs better than boxing alone for these systems. Yet, the quadratic nature of this curve can (just) be seen, and screening alone will become more expensive than boxing alone for extremely large systems. Not surprisingly, the fastest timings for large systems are when boxing and integral screening are both used.

Figure 3 shows the effect of altering ω for the same series of alkanes. Notice that the time saved




FIGURE 2. CPU times (s) for linear alkanes C_nH_{2n+2} with $\omega = 0.25$.




FIGURE 3. CPU times (s) for linear alkanes C_nH_{2n+2} using CASE + boxing + screening with various ω values.

by doubling ω decreases as ω increases and there is little to be gained by increasing ω beyond $1.0a_0^{-1}$ for these systems. This is because most of the integrals that remain with $\omega = 1.0a_0^{-1}$ involve diffuse basis functions and do not become negligible until very large ω values are used.

The effect of changing the integral accuracy for $C_{50}H_{102}$ is presented in Figure 4. The ratio of the conventional Coulomb CPU time to each other method's time reveals that all of these methods become relatively faster when the integral threshold is lowered. Although ONX with high accuracy is poor for these systems, its relative gain in speed as the accuracy is lowered is greatest, becoming more competitive with (but not faster than) CASE with boxing and screening (and $\omega = 0.25a_0^{-1}$).




FIGURE 4. Ratio of conventional Coulomb time to method time for $C_{50}H_{102}$ at various integral thresholds with $\omega = 0.25$.

Boxing and screening behave in similar ways with changes in the basis set (Fig. 5). The more uncontracted the basis set, the faster the method relative to the unattenuated calculation. This is a reflection of the time taken to test each *primitive* shell pair for its T_ω values. If a shell pair has only a small degree of contraction, the amount of testing is small. ONX shows similar behavior, except that it takes considerably longer when *d* functions are added to the basis set.

The performance of CASE on 2- and 3-dimensional systems can be seen in Figure 6. The solid curve shows timings for the conventional Coulomb method on a series of graphite sheets. Below this




FIGURE 5. Ratio of conventional Coulomb time to method time for $C_{50}H_{102}$ with various basis sets and $\omega = 0.25$.




FIGURE 6. CPU times (s) for graphite and diamond chunks with $\omega = 0.5$.

curve lie the boxing with screening and screening only curves. These CASE curves show linear behavior but are not as efficient as those for the linear alkanes. The CASE timings for the diamond chunks, however, are disappointing. It is only for $C_{84}H_{64}$ that CASE is faster and this is just the screening without boxing. The explanation for the poor performance lies in the physical size of the system. Our largest 3-dimensional structure ($C_{84}H_{64}$) is only 12 Å long on its largest side, compared with 19 Å for the largest graphite sheet used ($C_{96}H_{24}$) and 103 Å for the largest linear alkane ($C_{90}H_{182}$).

Concluding Remarks

We presented an efficient way to calculate the short-range Coulomb energy in molecules. This is the first $O(N)$ implementation of the CASE approximation, allowing rapid computation of CASE energies for large systems. A further speed-up was achieved by screening for negligible integrals. The increase in speed is heavily dependent on the shape of the system. For modest ω values, CASE represents a useful increase in speed even for systems of moderate size (e.g., $C_{50}H_{102}$) and, with the introduction of a correction for the long-range energy,^{18,19} CASE will offer a useful route to Hartree-Fock calculations of large molecules in the near future.

Acknowledgments

We would like to thank Terry Adams for several helpful discussions.

References

1. Greengard, L.; Rokhlin, V. *J Comput Phys* 1987, 73, 325.
2. White, C. A.; Johnson, B. G.; Gill, P. M. W.; Head-Gordon, M. *Chem Phys Lett* 1994, 230, 8.
3. Kutteh, R.; Apra, E.; Nichols, J. *Chem Phys Lett* 1995, 238, 173.
4. Strain, M. C.; Scuseria, G. E.; Frisch, M. J. *Science* 1996, 271, 51.
5. Dombroski, J. P.; Taylor, S. W.; Gill, P. M. W. *J Phys Chem* 1996, 100, 6272.
6. Perez-Jorda, J. M.; Yang, W. *J Chem Phys* 1997, 107, 1218.
7. Gill, P. M. W. *Adv Quantum Chem* 1994, 25, 141.
8. Adams, T. R.; Adamson, R. D.; Gill, P. M. W. *J Chem Phys* 1997, 107, 124.
9. Ewald, P. P. *Ann Phys* 1921, 64, 253.
10. Stoll, H.; Savin, A. In *Density Functional Methods in Chemistry*; Labanowski, J. K.; Andzelm, J., Eds.; Springer: New York, 1991; p 213.
11. Savin, A.; Flad, H.-J. *Int J Quantum Chem* 1995, 56, 327.
12. Adamson, R. D.; Dombroski, J. P.; Gill, P. M. W. *Chem Phys Lett* 1996, 254, 329.
13. Gill, P. M. W.; Adamson, R. D.; Pople, J. A. *Mol Phys* 1996, 88, 1005.
14. Gill, P. M. W.; Adamson, R. D. *Chem Phys Lett* 1996, 261, 105.
15. Adamson, R. D.; Gill, P. M. W. *J Mol Struct (Theochem)* 1997, 398, 45.
16. Gill, P. M. W. *Chem Phys Lett* 1997, 270, 193.
17. Lee, A. M.; Taylor, S. W.; Dombroski, J. P.; Gill, P. M. W. *Phys Rev A* 1997, 55, 3233.
18. Lee, A. M.; Adamson, R. D.; Gill, P. M. W. Unpublished manuscript.
19. Nair, N.; Lee, A. M.; Gill, P. M. W. Unpublished manuscript.
20. White, C. A.; Kong, J.; Maurice, D. R.; Adams, T. R.; Baker, J.; Challacombe, M.; Schwegler, E.; Dombroski, J. P.; Ochsenfeld, C.; Oumi, M.; Furlani, T. R.; Florian, J.; Adamson, R. D.; Nair, N.; Lee, A. M.; Ishikawa, N.; Graham, R. L.; Warshel, A.; Johnson, B. G.; Gill, P. M. W.; Head-Gordon, M. *Q-Chem, Version 1.2*; Q-Chem, Inc.: Pittsburgh, PA, 1998.

21. Gill, P. M. W.; Johnson, B. G.; Pople, J. A. *Int J Quantum Chem* 1991, 40, 745.
22. Lambert, C. G.; Darden, T. A.; Board, J. A., Jr. *J Comput Phys* 1996, 126, 274.
23. Rapaport, D. C. *Comput Phys Rep* 1998, 9, 1.
24. Burant, J. C.; Scuseria, G. E.; Frisch, M. J. *J Chem Phys* 1996, 105, 8969.
25. Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P. *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, 2nd ed.; Cambridge University Press: Cambridge, UK, 1992.
26. White, C. A.; Johnson, B. G.; Gill, P. M. W.; Head-Gordon, M. *Chem Phys Lett* 1996, 253, 268.
27. Johnson, B. G. Private communication.
28. Schwegler, E.; Challacombe, M. *J Chem Phys* 1996, 105, 2726.